arXiv:1601.07300v2 [cs.PL] 4 Feb 2016

# Recollection: an Alternative Restoration Technique for Constraint Programming Systems

Yong Lin and Martin Henz

School of Computing, National University of Singapore, Singapore
{linyong, henz}@comp.nus.edu.sg

**Abstract.** Search is a key service within constraint programming systems, and it demands the restoration of previously accessed states during the exploration of a search tree. Restoration proceeds either bottom-up within the tree to roll back previously performed operations using a trail, or top-down to redo them, starting from a previously stored state and using suitable information stored along the way. In this paper, we elucidate existing restoration techniques using a pair of abstract methods and employ them to present a new technique that we call recollection. The proposed technique stores the variables that were affected by constraint propagation during fix points reasoning steps, and it conducts neither operation roll-back nor recomputation, while consuming much less memory than storing previous visited states. We implemented this idea as a prototype within the Gecode solver. An empirical evaluation reveals that constraint problems with expensive propagation and frequent failures can benefit from recollection with respect to runtime at the expense of a marginal increase in memory consumption, comparing with the most competitive variant of recomputation.

## 1 Introduction

Constraint programming is a tool for modelling and solving constraint satisfaction problems. To support this tool, constraint programming systems provide a set of services to conduct constraint-based search. Among these services, restoration implements the process of restoring a previously reached computation state. A straightforward albeit memory-intensive form of restoration is afforded by making a *copy* of each previously visited state; restoration then just retrieves the expected one from memory. Less memory-intensive alternatives construct the required state based on suitable *chunks* of information collected during search. This reconstruction can be achieved by either *trailing* or *recomputation*.

Implementations of logic programming (Prolog), an ancestor of constraint programming, use trailing as the main restoration technique. Most of today's CP systems are constraint logic programming [8] systems that evolved from Prolog and inherited its trailing-based restoration paradigm, such as ECL$^i$PS$^e$ [1], CHIP [7], clp(FD) [6] and SICStus Prolog [4]. This technique rolls back previously performed operations to restore an earlier state by utilizing undo information maintained in a trail structure, proceeding in a *bottom-up* direction within

the search tree. By contrast, recomputation restores a state by redoing previous computation, which follows a *top-down* manner. Recomputation, combined with copying in various ways, is the restoration technique underlying the concurrent constraint programming system Mozart/Oz [14], as well as the constraint programming solver Gecode [16].

This paper investigates a top-down restoration technique that we call *recollection*, which avoids redoing computation, while consuming much less memory than copying. Its central idea is simple; we incrementally memoize the variable domains that were updated to reason fix point state for later restoration.

*Plan of the paper.* The following Section 2 introduces some basic notions referred to in the rest of the paper. Section 3 briefly reviews the current available restoration techniques using a pair of abstract methods, and Section 4 presents our proposed *recollection* technique together with its implementation issues. Section 5 describes the evaluation of our implemented recollection on selected benchmark problems, and we conclude in Section 6.

## 2 Constraint-based Problem Solving

This section introduces the basic terms and reviews the main principles of constraint programming to provide the context for the rest part of this paper.

### 2.1 Basics

A constraint satisfaction problem defines a finite set of variables together with constraints. Each variable has a domain to specify the finite set of values it can take, and a variable is *fixed* if its domain is a singleton set. The conjunction of variables forms a *store*, which maps variables to their domains. A constraint describes a certain relationship over a subset of variables to restrict their value combinations that can appear in the solutions to the problem. In constraint programming systems, a constraint is implemented as a *propagator*, and propagators amplify their store by executing their filtering algorithms to rule out inconsistent values from variable domains.

A store and its connected propagators form a *state*, and the store provides a communication channel for its propagators. Specifically, variable domain changes by running a propagator will be reflected in the store, which may trigger the execution of other propagators to filter more values from other variable domains. This process is called *constraint propagation* (propagation for short). A propagation is *strong* if it is able to change the majority of variables within the problem; otherwise, the propagation is *weak*. During propagation, if any variable's domain becomes empty, the engine infers an *inconsistency* and the current state becomes *failed*; if all variables become fixed, the state is in a *solved* status and represents a *solution*. Nevertheless, constraint propagation alone is generally insufficient to identify a solution; it will reach a *fix point* when constraint propagation cannot make further changes to the store.

In such a circumstance, search is required and it equivalently splits the fix point state into multiple further constrained states. The split is called *branching* and it is achieved by invoking a *brancher* in a constraint programming system. The invocation of a brancher generates a *choice* of the fix point state, and the choice includes multiple mutually exclusive constraint alternatives[1]. The search process *commits* an alternative of the choice to visit the corresponding further constrained state, and constraint propagation can resume at the further constrained state. A choice is *open* if it has an uncommitted alternative; otherwise, it is *closed*.

### 2.2 Search

Search is a central service in constraint programming systems and it is programmed as a *search engine*. The search consists of *exploration* and *restoration*. Exploration alternates between constraint propagation and branching, and this leads to a tree of states, the *search tree*. In such a search tree, branches represent constraints, internal nodes are fix points and the leaf nodes are either solved or inconsistent states. Inconsistency indicates a false search direction, and restoration needs to recover a previously accessed internal state, *target state*, to guide exploration to other part of the search tree. An important step for restoration is to decide which state to restore after encountering an inconsistency, and in this paper, we focus on *chronological backtracking*.

In the search tree, the *root* state is the fix point state reasoned on the original problem, and the *current* state is the one that search engine is interacting. The branches and fix points between the *root* and *current* node form a *path*[2]. In a constraint programming system, the path is designed to store the information that has been utilized to reach current state from root, and the system makes use of this information to conduct restoration.

Algorithm 1 describes a DFS search without committing to a particular restoration technique. In this pseudo-code, the variable $S$ refers current state; the path information is maintained explicitly using a stack $ST$. The method *Propagate*() (Line 3) conducts constraint propagation within state $S$ and the *log* records the changes enforced by the propagation to state $S$. A switch statement responds according to the propagation result. Specifically, if a state is solved, it will be returned as a solution. If propagation exhibits an inconsistency, the search engine will call the method *Restore* to restore a state, from which to explore the second branch of the node. If a fix point is reasoned, the search engine will call Branch() to generate a choice of the state and then commits it to the first alternative by the Commit() method; subsequently, a chunk will be constructed (Line 18) and pushed onto the stack $ST$.

In the pseudo-code, the methods *Record* and *Restore* form a pair of abstract methods, whose implementations determine the way to conduct restoration. In the following sections, we discuss restoration techniques by describing this pair of methods.

[1] We restrict our discussion to binary choice in this work.

[2] We focus on sequential search in this paper and therefore there is a single path.

**Algorithm 1** Depth First Search

```
Input: State S, Stack ST
Output: Solution State
 1: while true do
 2:    Log log
 3:    switch (Propagate(S, log))
 4:    case solved:
 5:       return S
 6:    case inconsistency:
 7:       S ← Restore(S, ST, log)
 8:       if S = NULL then
 9:          return non-solvable
10:       end if
11:       Chunk chunk ← getTop(ST)
12:       Choice choice ← getChoice(chunk)
13:       Commit(S, choice, second)
14:       break
15:    case fix_point:
16:       Choice choice ← Branch(S)
17:       Commit(S, choice, first)
18:       Chunk chunk ← Record(S, choice, log)
19:       Push(ST, chunk)
20:       log ← ∅
21:    end switch
22: end while
```

## 3 Restoration

The restoration of a previously accessed state can be achieved by either reconstruction or memorization. Memorization has been implemented by copying (Section 3.2); reconstruction can be approached by trailing (Section 3.1) and recomputation (Section 3.3). In this section, we briefly go through these restoration techniques together with their characteristics.

### 3.1 Trailing

Trailing-based systems implement the *Record* method to accumulate operation undo information in a trail structure (it is the stack $ST$ in this context). Conceptually, the undo information is expected to describe the changes enforced to the state. In practical implementations, systems mostly prefer maintaining the original images before updates. Examples are single-value [3], Time-Stamping and Multiple-Value trail (see [2]) etc. For a comprehensive description, please refer to [9].

The updating of a trail structure interleaves with propagation since variables keep changing during propagation. In Algorithm 1, a global data structure *log* is introduced to collect undo information. If propagation leads to a fix point, the collected undo information will be wrapped into the chunk and pushed onto $ST$. Algorithm 2 shows that trailing-based restoration first cancels the operations stored in the *log*. Subsequently, the chunks in $ST$ are accessed one by one to conduct a step-wise roll-back; this process iterates until it backtracks to a state that owns an open choice.

**Algorithm 2** *Restore*$_{trail}$

```
Input: State S, Stack ST, Log log
Output: State S
 1: undo(S, log)
 2: Chunk chunk ← getTop(ST)
 3: Choice choice ← getChoice(chunk)
 4: while choice has no uncommitted alternative do
 5:     log ← getLog(chunk)
 6:     undo(S, log)
 7:     Pop(ST)
 8:     if Size(ST) = 0 then
 9:         return NULL
10:     end if
11:     chunk ← getTop(ST)
12:     choice ← getChoice(chunk)
13: end while
14: return S
```

**Algorithm 3** EXPOSE_OPEN_CHOICE

```
Input: Stack ST
Output: Stack ST
 1: Chunk chunk ← getTop(ST)
 2: Choice choice ← getChoice(chunk)
 3: while choice has no uncommitted alternative do
 4:     Pop(ST)
 5:     if Size(ST) = 0 then
 6:         return NULL
 7:     end if
 8:     chunk ← getTop(ST)
 9:     choice ← getChoice(chunk)
10: end while
11: return ST
```

### 3.2 Copying

Copying-based state restoration defines the *Record*$_{copy}$ method to store a copy of the entire state in each constructed chunk. The corresponding *Restore*$_{copy}$ method is straightforward: retrieve the chunk containing the expected state and return. For later reuse, the function EXPOSE_OPEN_CHOICE in Algorithm 3 implements chronological backtracking by popping chunks from the stack until it recognizes a chunk that contains an open choice.

**Algorithm 4** *Restore*$_{copy}$

```
Input: State S, Stack ST, Log log (ignored)
Output: State S
1: delete S
2: call EXPOSE_OPEN_CHOICE
3: Chunk chunk ← getTop(ST)
4: return S ← getState(chunk)
```

Copying can be more memory intensive than trailing, and the intensive memory consumption may introduce non-negligible garbage collection cost, especially for large problems with a substantial number of variables; Schulte has conducted

a deep study on this topic in [13]. Nevertheless, it is essential to ensure multiple states are simultaneously available for parallel search, and the concurrent constraint programming system Mozart/Oz [15,14] adopts copying as one of its restoration techniques.

### 3.3 Recomputation

Recomputation implements $Record_{recomp}$ to store the constraints that were generated at fix points states and the $Restore_{recomp}$ computes from root state downwards to restore the target state, using the set of the introduced constraints between them. Algorithm 5 depicts the pseudcode for implementing recomputation. In this depiction, the **for** loop commits constraints in a batch [5] rather than step-wisely as implemented in Gecode.

**Algorithm 5** $Restore_{recomp}$

**Input:** State $S$, Stack $ST$, Log $log$ (ignored)
**Output:** State $S$
1: **delete** $S$
2: $S \leftarrow$ getRootState($ST$)
3: **call** EXPOSE_OPEN_CHOICE
4: **for each** $chunk \in ST$ **do**
5: $choice \leftarrow$ getChoice($chunk$)
6: Commit($S$, $choice$, $oldAlternative$)
7: **end for**
8: **return** $S$

Recomputation consumes little memory at the expense of a considerable runtime penalty since computation from scratch is performed whenever a failure occurs. To alleviate, a hybrid scheme can be formed by placing state copies somewhere within the search tree; typical examples are *fixed recomputation* and *adaptive recomputation*. Fixed recomputation places a state copy after every $d$ exploration steps, where $d$ is a constant called *copying distance*. Fixed recomputation was extended to *adaptive recomputation*: when recomputing from $S_1$ to $S_2$, it will put an additional state copy in the middle between $S_1$ and $S_2$ to shorten future recomputation distance. Adaptive recomputation has been demonstrated as a competitive restoration technique [13] and is supported by the Mozart/Oz as well as Gecode systems.

## 4 Recollection

As we have reviewed in the previous section, recomputation needs to re-execute the propagators' filtering algorithm to compute a visited state, while copying stores every reasoned fix point state. Unlike copying, trailing instead is concerned with data structures within states and records the information to roll back performed changes when needed. In this section, we propose an alternative restoration technique that we call *recollection*, which memoizes the propagation

updated variable domains to reach a fix point for achieving restoration. By following the presentation paradigm in the previous section, we intend to explain the idea of recollection by describing the implementation of the pair of abstract methods, *Record* and *Restore*.

### 4.1 The Record Method.

As illustrated in Algorithm 1, the *Record* method is invoked after a fix point has been reached. Since the goal of recollection is to memoize propagation affected variable domains for restoration, we should identify the set of variables that were changed during propagation and memorize their domains, as explained in the **for** loop in Algorithm 6.

**Algorithm 6** $Record_{recollect}$

```
Input: State S, Choice choice, Log log (ignored)
Output: Chunk chunk
1: Domain doms ← ∅
2: for each var ∈ Variables(S) do
3:     if isChanged(var) then
4:         doms ← doms ∪ recordDomain(var)
5:     end if
6: end for
7: return Chunk(choice, doms)
```

### 4.2 The Restore Method.

To restore a state, recollection faces an issue that the required variable domains scatter over the chunks above the target node within the search tree. Therefore, recollection should collect variable domains efficiently for restoration, and we propose two manners: the variable-centered and the chunk-centered.

*Variable-centered restoration.* The variable-centered approach, showing in Algorithm 7, picks one variable *var* at a time and searches the stack $ST$ in a top-down direction (moving in the search tree in a bottom up direction!) for the first chunk that contains its domain (Line 7 to 10) and then reconstructs it (Line 11). Suppose the constraint problem imposes $M$ variables and the number of chunks within the stack is $N$. In the worst case, this approach would conduct $N \times M$ chunk access operations to achieve a restoration; a rather weak propagation problem may approach such a worst-case scenario. On the other hand, most variables can retrieve their domains at the top chunk in the presence of a strong propagation problem and thus the restoration would conduct slightly more than $M$ times chunk access operations.

**Algorithm 7** *Restore*$_{variable\text{-}centered}$

```
Input: State S, Stack ST, Log log (ignored)
Output: State S'
 1: delete S
 2: S ← getRootState(ST)
 3: call EXPOSE_OPEN_CHOICE
 4: for each var ∈ Variables(S) do
 5:    Integer index ← Size(ST) - 1
 6:    chunk ← getChunk(ST, index)   /*scan from stack top*/
 7:    while Domain(var) ∉ chunk do
 8:       index ← index - 1    /*move to next chunk location*/
 9:       chunk ← getChunk(ST, index)
10:    end while
11:    Reconstruct(var, chunk)
12: end for
13: return S
```

**Algorithm 8** *Restore*$_{chunk-centered}$

```
Input: State S, Stack ST, Log log (ignored)
Output: State S'
 1: delete S
 2: S ← getRootState(ST)
 3: call EXPOSE_OPEN_CHOICE
 4: Integer index ← Size(ST) - 1
 5: while index ≥ 0 do
 6:    Chunk chunk ← getChunk(ST, index)
 7:    for each var ∈ Variables(chunk) do
 8:       if var has not been reconstructed then
 9:          Reconstruct(var, chunk)
10:       end if
11:    end for
12:    index ← index - 1
13: end while
14: return S
```

*Chunk-centered restoration.* By contrast, the chunk-centered approach, showing in Algorithm 8, scans $ST$ in a top-down manner (moving bottom-up in the search tree) and keeps track of reconstructed domains. Within each chunk, all memoized variables are scanned and a variable domain will be reconstructed if it has not been reconstructed yet (Lines 6–14). This query scheme requires to access the stack once during a restoration, regardless of whether the problem exhibits weak or strong propagation, and it has a slight runtime advantage over variable-centered restoration according to our experimental results. The experiments of the next section have been conducted using the chunk-centered restoration; we introduce an index to facilitate the recognition of reconstructed variables.

### 4.3 Variants

Our discussion on recollection so far assumes that a single state is maintained at the search tree root and thus restoration begins from scratch. We observer that recollection alone usually incurs a significant runtime penalty, especially for the problems that create a deep search tree. This issue attributes to the process of searching chunks when recollection intends to restore a state. To alleviate the penalty, we extend recollection to the variants of *fixed recollection* and *adaptive*

*recollection*, analogous to the recomputation variants; they place state copies somewhere in the search in the way that have been explained in Section 3.3.

A recent studied restoration technique in [12] intends to combine trailing, copying and recomputation to construct a hybrid architecture. In addition to variable domains, this hybrid scheme is also concerned about propagators and propagator states as well as the dependencies between variables and propagators. On the other hand, recollection concentrates on variable domains and is transparent with other detail. Meanwhile, recollection also can be orthogonal with the implementation of search.

### 4.4 Implementation Issues

Gecode is a system that develops computation space (space for short) as its central concept. A computation space encapsulates computations and is home to variables, propagators and branchers. Computation spaces are first-class structures that can be copied and programmed to construct search engines using its provided operations, a computation space corresponds a node in the search tree. In this sub-section, we discuss three key issues to implement recollection within Gecode system.

**Memory Management** The implementation of method $Record_{recollect}$ consumes memory, but Gecode's *memory manager* is centered on spaces, while chunks live most naturally outside spaces. Hence, we should provide a memory management policy to support the memorized variable domains. In our prototype, we investigated two options to address the memory management issue: (1). allocate memory dynamically for recorded each variable; (2). calculate the expected total memory size first and then allocate once. The allocated memory is freed when its chunk is popped from the stack. Our experiments reveal that approach (1) is marginally more runtime efficient for problem with weak propagation, while approach (2) is more suitable for problems of many variables with strong propagation. In our experimental prototype [11], we can switch between the two alternatives by setting a compile-time flag. Note that we use approach (1) in all experiments reported in the next section.

**Domain Change Detection** A key implementation challenge of the method $Record_{recollect}$ in Section 4.1 is identifying the set of variables that were updated to reason each fix point. Fortunately, our chosen recollection implementation platform, Gecode, developed an abstraction called *advisor* [10] to optimize constraint propagation. We make use of the instantiated advisor structures to recognize the changed variables. For detail of the technique to achieve this goal, please refer to the source code of our published recollection implementation [11].

**Variable Reconstruction** For a variable domain, the removal of values by constraint propagation may break it into multiple intervals. In Gecode, each

interval is implemented as a *range* data structure and multiple ranges within a variable implementation will be chained as a list. In addition to updating the values represented on each range, recollection may be required to adjust the current chain length as well. We developed two methods to achieve this goal.

The first option simply destroys the current chain and then rebuilds another one. This method is straightforward to implement, but frequent chain destroying may introduce intensive garbage collection for the system, especially when dealing with long chains. Therefore, this alternative should perform better to deal with short and medium length chains. By contrast, the second approach either trims or extends the current chain to the new length. In our prototype, we have implemented both alternatives and they are configurable through a compile-time flag. For the experiments in evaluation section, we employ the second approach.

## 5 Evaluation

This section empirically evaluates recollection over a set of benchmark problems: Section 5.1 describes the overall evaluation setting; Section 5.2 compares recomputation with recollection and Section 5.3 intends to extend comparison to other restoration techniques.

### 5.1 Configurations

We used a PC system that is equipped with Intel Core 2 Quad processor Q9550, running the Ubuntu operating system 12.04 in a 32-bit mode with four Gigabyte main memory. We built our prototype [11] on top of Gecode version 3.7.3 [16], which also served as the reference instance for comparison; the source code was compiled by G++ version 4.6.3. Each collected runtime[3] value is an arithmetic mean of 20 runs with a variation coefficient less than 2%; memory measurements is the *peak* amount of the memory occupation.

As benchmarks, we used finite domain integer and Boolean problems. They were selected to cover multiple constraints, spawn a varying number of propagators and impose different propagation intensity. Meanwhile, they cover first, all and best (branch-and-bound) solution search. We limit ourselves to the problems included in the Gecode repository and stick to the configuration of propagation consistency level and branching strategy of the original scripts.

The set of selected benchmark problems are: the Queens problem modelled by either a quadratic number of disequality constraints or three global constraints that generalize *all-different*; the magic-square puzzle of size 5; a sport league problem with 22 teams; the black hole patience game; Balanced Incomplete Block Design (BIBD), the knights tour problem of size 22; the Pentominoes problem; the Alpha crypto-arithmetic puzzle; the Langford's number problem with 3 by 9 values and; Golomb-Ruler problem of size 10 and the problem of independent sets in graph (Ind-Set). Table 1 lists the characteristics of these problems, where

[3] We take wall clock time in this work.

| Problem | Sols | Propagators | Propagations | Nodes | Failures | Depth |
|---|---|---|---|---|---|---|
| Queens(100) | one | 14,850 | 16,821 | 138 | 22 | 96 |
| Queens-S(100) | one | 3 | 428 | 138 | 22 | 96 |
| Magic-Square(5) | one | 15 | 2,292,251 | 144,471 | 72,227 | 33 |
| Sport-League(22) | one | 1,199 | 207,066 | 2,273 | 1,035 | 249 |
| Black-Hole | one | 742 | 986,542 | 5,284 | 2,631 | 47 |
| BIBD | one | 9,693 | 912,464 | 2,625 | 1,306 | 968 |
| Knight(22) | one | 1 | 74,610 | 40,184 | 19,877 | 451 |
| Pentominoes | one | 81 | 6998 | 143 | 64 | 27 |
| Alpha | all | 21 | 136,179 | 14,871 | 7,435 | 49 |
| Langford-Num | all | 37 | 22243 | 303 | 149 | 17 |
| Golomb-Ruler(10) | optimal | 39 | 2,760,799 | 39,875 | 19,928 | 33 |
| Ind-Set | optimal | 21 | 101,317 | 29,849 | 14,895 | 40 |

**Table 1.** Characteristic of Benchmark Problem Search Trees

the *propagations* are the numbers collected when using adaptive recomputation for restoration with default argument settings. For the original scripts, refer to the Gecode distribution in [16].

### 5.2 Recomputation and Recollection

The proposal of recollection was motivated by recomputation in a sense of skipping the re-execution of constraint propagation, hence the foremost performance comparison should be with recomputation. As explained in Section 3, recomputation can derive variants by combining with copying, and adaptive recomputation generally exhibits superior runtime performance compared with other recomputation schemes [13]. Similarly, adaptive recollection exposes the most runtime competitive recollection scheme. We therefore first focus on a direct comparison between adaptive recomputation and adaptive recollection, fixing the copying distance to *eight* in both cases.

Table 2 depicts the experimental results, which demonstrates that neither recomputation nor recollection can impose a consistent performance advantage over the other for solving all problems. Recollection hardly improves the runtime of the problems with shallow search trees and limited number of failures such as Pentomonies and Langford-Number; or even leads to an inferior runtime, as in Alpha and Magic Squares. Nevertheless, recollection can be runtime competitive for finite domain integer problems with deep search trees and intensive search failures such as Sport-League, Golomb-Ruler and Knights. On these problems, recollection is able to make an runtime improvement by investing a small amount of additional memory than adaptive recomputation.

Boolean problems can hardly benefit from recollection, even though a Boolean problem would explore a rather deep search tree (BIBD) and incur intensive failures (Ind-Set). This is mainly because that a Boolean variable contains at most

| Problems | Recomputation | | Recollection | |
|---|---|---|---|---|
| | Time(ms) | Mem(KB) | Time (ms) | Mem(KB) |
| Queens(100) | 16 | 4,301 | 15 | 4,663 |
| Queens-S(100) | 1 | 240 | 2 | 602 |
| Magic-Square(5) | 579 | 63 | 653 | 73 |
| Sport-League(22) | 352 | 7,710 | 331 | 7,937 |
| Black-Hole | 535 | 1,927 | 508 | 1998 |
| BIBD | 573 | 4,678 | 575 | 4784 |
| Knights(22) | 1,858 | 4,460 | 1,704 | 4,592 |
| Pentominoes | 20 | 1,158 | 19 | 1,173 |
| Alpha | 55 | 45 | 66 | 50 |
| Langford-Number | 13 | 132 | 13 | 135 |
| Golomb-Ruler(10) | 556 | 69 | 547 | 70 |
| Ind-Set | 58 | 41 | 68 | 43 |

**Table 2.** Comparison Adaptive Recomputation and Adaptive Recollection

two values; recollection memoized Boolean domains are not dense enough to compete with the recomputation via re-runing propagation algorithms.

The measurements in Table 2 was conducted with a specific copying distance (*eight*), and this may impose a concern that the fixed copying distance may conclude skewed results. To dispel this concern, we ran Sport-League and Knights problem in adaptive recomputation and adaptive recollection respectively over a range of copying distance values, and Table 3 illustrates the collected runtime measurements.

| | Copying Distance ($d$) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | $d=1$ | $d=3$ | $d=5$ | $d=10$ | $d=20$ | $d=40$ | $d=80$ | $d=160$ | $d=320$ |
| SportLeague(*recomp*) | 337 | 330 | 341 | 350 | 351 | 355 | 359 | 360 | 359 |
| SportLeague(*recoll*) | 336 | 325 | 326 | 330 | 329 | 333 | 334 | 336 | 335 |
| **Time** $\Delta$(ms) | 1 | 5 | 15 | 20 | 22 | 22 | 25 | 24 | 24 |
| Knights(*recomp*) | 1598 | 1737 | 1830 | 1856 | 1855 | 1868 | 1872 | 1855 | 1864 |
| Knights(*recoll*) | 1589 | 1646 | 1695 | 1712 | 1711 | 1703 | 1697 | 1700 | 1714 |
| **Time** $\Delta$(ms) | 9 | 91 | 135 | 144 | 144 | 165 | 175 | 155 | 150 |

**Table 3.** Sport-League and Knight run over a range of copying distances

Table 3 shows that adaptive recollection is able to adjust quickly to converge to a small runtime interval, even though the copying distance is set to a large value. This observation indicates that the setting of copying distance is not as significant as one may have imagined, which confirms and generalizes the corresponding original observation reported on adaptive recomputation in [13]. As

observed from the table, the runtime difference between recomputation and recollection initially grows as copying distance increases and then shrinks somewhat after reaching a peak performance gap (at $d$=80 in both cases); afterwards, it stays almost stable with the further increase of the copying distance.

### 5.3 Extended Comparison

In addition to recomputation, the evaluation would be more thorough if it includes copying and trailing. Fortunately, copying-based restoration can be easily obtained by setting the copying distance to *one* in Gecode. As for trailing, it is absent in current Gecode system. To implement a trailing-based restoration, we are expected to clear many techniques choices such as trail structure, and these factors are key to the performance of a trailing-based system. Meanwhile, Gecode centers on copying with recomputation and it has developed intensive techniques to optimize its underlying restoration. In such a circumstance, a concern about fair comparison would come out. Alternatively, another way is to employ a set of trailing-based systems to carry out platform-crossing empirical evaluation. Such a work would indeed be a significant contribution, and Schulte has conducted a similar investigation in [13]. However, this exceeds our resource in this paper, the goal of which is to propose an alternative restoration technique. Instead, we take the systematic platform-crossing study as one of our main future research works. Nevertheless, we intend to illuminate the comparison with other restoration techniques as illustrated in Table 4.

| | Copying | | Recomputation | | Recollection | |
|---|---|---|---|---|---|---|
| Problems | Time(ms) | Mem(KB) | Time(ms) | Mem(KB) | Time(ms) | Mem(KB) |
| Queens(100) | 39 | 26076 | 16 | 4301 | 15 | 4663 |
| Queens-S(100) | 2 | 1662 | 1 | 240 | 2 | 602 |
| Queens(200) | 3837 | 170669 | 4298 | 6224 | 4626 | 9027 |
| Queens-S(200) | 1996 | 8560 | 2171 | 1066 | 2479 | 3866 |
| Alpha | 51 | 54 | 55 | 45 | 66 | 50 |
| Magic-Square | 487 | 100 | 579 | 58 | 652 | 67 |
| Knights(18) | 31 | 11271 | 20 | 1596 | 19 | 1681 |
| Knights(22) | 1598 | 30159 | 1858 | 4460 | 1704 | 4592 |
| Golomber-Ruler | 469 | 77 | 550 | 61 | 550 | 63 |

**Table 4.** Comparison with other restoration techniques

The table reveals that recollection can consume much less memory than copying as one can expect for most problems. For runtime, recollection generally cannot compete with copying except on few cases. In [12], it implemented a simplified trailing in Gecode to compare with other restoration strategies. Its conducted evaluation demonstrates that copying is slightly preferable if problems impose strong propagation, while trailing benefits a lot in weak propagation

problems. This finding may entail a promising comparison between recollection and trailing since recollection has developed methods (variable-center restoration in Section 4.2) to aim for efficient restoration in strong propagation problems.

## 6 Conclusion

We proposed a new top-down state restoration technique that we call *recollection* and present it in an exposition using the abstract methods *Record* and *Restore*. Unlike copying, recollection exhibits a finer granularity and only memoizes the variable domains that were changed during each fix point computation; compared with recomputation, it would not re-run the propagators' built-in filtering algorithms for restoration; as opposed to trailing, recollection approaches in a top-down direction within the search tree, and it is orthogonal with search.

Our experimental evaluation reveals that recollection has the opportunity to improve runtime against adaptive recomputation on integer problems with deep search trees and intensive search failures, at the expense of moderate memory investment. Our current prototype was built on Gecode system and has investigated several implementation alternatives to develop this idea, and we cautiously conjecture more significant performance improvements if one takes further efforts to explore more aggressive techniques to optimize recollection implementation. We hope that the presented recollection technique is significant enough to be added to the set of tools available to designers and implementers of constraint programming systems. Exciting future research topics can be exploring parallel and non-chronological backtracking search by utilizing recollection as the underlying restoration strategy. Additionally, an in-depth comprehensive system-crossing performance study on various restoration techniques would be another significant contribution.

### Acknowledgements

We are very grateful to Christian Schulte and Guido Tack for unfailingly resourceful answers to our pestering queries and their (cautious) words of encouragement.

## References

1. Abderrahamane Aggoun, David Chan, Abderrahamane Aggoun, David Chan, Pierre Dufresne, Eamon Falvey, Alexander Herold, Geoffrey Macartney, Micha Meier, David Miller, Shyam Mudambi, Bruno Perez, Emmanuel Van Rossum, Joachim Schimpf, and Periklis Andreas Tsahageas. Eclipse 3.5 user manual, 1995.
2. Abderrahmane Aggoun and Nicolas Beldiceanu. Constraint logic programming. chapter Overview of the CHIP compiler system, pages 421–435. MIT Press, Cambridge, MA, USA, 1993.
3. Hassan Aït-Kaci. *Warren's Abstract Machine: A Tutorial Reconstruction*. Logic Programming Series. The MIT Press, Cambridge, MA, USA, 1991.

4. Mats Carlsson and Per Mildner. SICStus Prolog—the first 25 years. *TPLP*, 12(1–2):35–66, 2012.
5. Chiu Wo Choi, Martin Henz, and Ka Boon Ng. Components for state restoration in tree search. In *Proceedings of the 7th International Conference on Principles and Practice of Constraint Programming*, CP '01, pages 240–255, London, UK, 2001. Springer-Verlag.
6. Philippe Codognet and Daniel Diaz. Compiling constraints in clp(fd). *The Journal of Logic Programming*, 27(3):185–226, 1996.
7. Mehmet Dincbas, Pascal Van Hentenryck, Helmut Simonis, Abderrahmane Aggoun, Thomas Graf, and Françoise Berthier. The constraint logic programming language chip. In *Proceeding of the International Conference on Fifth Generation Computer Science FGCS-88*, pages 693–702, 1988.
8. Joxan Jaffar and J.-L. Lassez. Constraint logic programming. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '87, pages 111–119, New York, NY, USA, 1987. ACM.
9. Joxan Jaffar and Michael J. Maher. Constraint logic programming: a survey. *The Journal of Logic Programming*, pages 503 – 581, 1994. Special Issue: Ten Years of Logic Programming.
10. Mikael Z. Lagerkvist and Christian Schulte. Advisors for incremental propagation. In Christian Bessiere, editor, *Thirteenth International Conference on Principles and Practice of Constraint Programming*, volume 4741 of *Lecture Notes in Computer Science*, pages 409–422, Providence, RI, USA, September 2007. Springer-Verlag.
11. Yong Lin. Prototypical implementation of recollection based on Gecode. `www.comp.nus.edu.sg/~henz/recollection`, May 2013.
12. Raphael M. Reischuk, Christian Schulte, Peter J. Stuckey, and Guido Tack. Maintaining state in propagation solvers. In Ian Gent, editor, *Fifteenth International Conference on Principles and Practice of Constraint Programming*, volume 5732 of *Lecture Notes in Computer Science*, pages 692–706, Lisbon, Portugal, September 2009. Springer-Verlag.
13. Christian Schulte. Comparing trailing and copying for constraint programming. In *Proceedings of the 1999 International Conference on Logic Programming*, pages 275–289, Cambridge, MA, USA, 1999. Massachusetts Institute of Technology.
14. Gert Smolka. The Oz programming model. In Jan van Leeuwen, editor, *Lecture Notes in Computer Science*, volume 1000, pages 324–343, Berlin, 1995. Springer-Verlag.
15. Gert Smolka, Martin Henz, and Jörg Würtz. Object-oriented concurrent programming in Oz. In Vijay Saraswat and Pascal Van Hentenryck, editors, *Principle and Practice of Constraint Programming*, pages 29–48, Cambridge, MA, USA, 1995. MIT Press.
16. Gecode Project Team. Generic COnstraint Development Environment. `www.gecode.org`, June 2012.